

REVIEW

Constraints on Learning as Default Assumptions: Comments on Merriman and Bowman's "The Mutual Exclusivity Bias in Children's Word Learning"

AMANDA L. WOODWARD AND ELLEN M. MARKMAN

Stanford University

The Mutual Exclusivity Bias in Children's Word Learning. By W. E. MERRIMAN AND L. L. BOWMAN. In *Monographs of the Society for Research in Child Development*, 54, Serial No. 220, 1989.

Merriman and Bowman (1989) explore one proposed constraint on word learning, mutual exclusivity, the assumption that each object has only one label. Chief among their contributions is treatment of mutual exclusivity as a default assumption—a probabilistic bias which can be overridden. We elaborate this view to address misconceptions of the notions of biological constraints that pervade recent discussions of constraints on word learning. Misconstruing constraints as rigid, absolute responses instead of probabilistic biases has led researchers to interpret any violation as invalidating a given constraint. More confusion surrounds questions about the origins of the constraints. We dispute the idea that the age of appearance of lexical constraints reveals whether the constraints are innate, and argue that current discussions of the innateness of constraints are oversimplified. In this case, we also question the appropriateness of Merriman and Bowman's methodology for use with 2-year-olds and challenge their conclusion that mutual exclusivity is absent in children under 2½. Merriman and Bowman's thoughtful conceptual analysis establishes several distinct ways in which mutual exclusivity can be manifested. Thus, putative counterexamples occur when an investigator tests for only one consequence of mutual exclusivity and ignores its other possible implications. Merriman and Bowman's studies with children from 2½ on document that each of these alternative ways of preserving mutual exclusivity guides word learning. © 1991 Academic Press, Inc.

A fundamental problem for language learners is figuring out what new words mean. Several researchers (Carey, in press; Clark, 1988; Dockrell & Campbell, 1986; Golinkoff, Hirsh-Pasek, Baduini, & Lavalley, 1985; Hutchinson, 1984, 1986; Markman, 1987, 1989, 1990, in press; Markman & Hutchinson, 1984; Markman & Wachtel, 1988; Mervis, 1989; Waxman, 1989) have argued that children are equipped with early biases or con-

This work was supported by an NSF graduate fellowship to Woodward and NIH Grant HD20382 to Markman. We are grateful to Anne Fernald and Dare Baldwin for their helpful comments on an earlier version of the manuscript.

straints that guide their hypotheses about word meanings. The notion of constraints on language learning is controversial: while some researchers reject the idea entirely (e.g., Nelson, 1988), those who accept it disagree as to whether children could in principle acquire language without such constraints. MacWhinney (1989), for example, treats constraints as heuristics learned by the child which are useful rather than essential. On the other hand, others argue that some such biases are essential to early word learning (Clark, 1983, 1987, 1988; Gelman, 1990a, 1990b; Golinkoff et al., 1985; Markman, 1987, 1989, 1990 in press; and Waxman, 1989).

In their recent monograph, Merriman and Bowman (1989) explore one proposed constraint on word learning, *mutual exclusivity*—the assumption that each object has only one label. In our discussion of Merriman and Bowman's work we hope to accomplish several goals: (1) to address misconceptions of the notions of biological constraints that pervade recent discussions of mutual exclusivity and other proposed constraints on word learning; (2) to argue that a child's lexicon can be a poor source of evidence about the process by which words were acquired; (3) to evaluate an alternative hypothesis which Merriman and Bowman propose accounts for some of the evidence in support of mutual exclusivity; (4) to dispute the idea that the age of appearance of lexical constraints reveals whether the constraints are innate and to argue that recent discussions of the origins of constraints oversimplify the issues; (5) to question the appropriateness of Merriman and Bowman's methodology for use with 2-year-olds; and, (6) to challenge their conclusion that mutual exclusivity is absent in children under 2½ years of age. Before discussing these issues, we will briefly summarize the perspective that provides the background to Merriman and Bowman's work.

WORD LEARNING AS AN INDUCTIVE PROBLEM

Word learning is an inductive feat (Quine, 1960). One well-known problem of induction is that the evidence is always too impoverished to logically eliminate all but one hypothesis (Goodman, 1955; Quine, 1960). Quite the contrary, for word learning an infinite number of hypotheses are consistent with information obtained through ostensive definition. When one points to an object and labels it, "see the dog," one is simultaneously pointing to its color, size, material, position, attractiveness, value, and an infinite set of other properties; for example, it weighs less than 100 pounds, less than 101 pounds, and so on. Thus, word learning presents a problem of induction that must somehow be solved by very young children, who have known limitations on their information processing abilities. Ethology offers insights as to what might contribute to a solution to

the inductive problem that word learning poses, namely, that hypotheses can be constrained in such a way as to greatly facilitate learning.

It is hard to see how children could acquire language as rapidly as they do without biases that enable them to rule out many alternative hypotheses for the meaning of a word and that lead them instead to focus on hypotheses that are reasonably likely to be correct (Carey, in press; Clark, 1987, 1988; Golinkoff et al., 1985; Hutchinson, 1984, 1986; Markman, 1987; 1989, 1990, in press; Markman & Hutchinson, 1984; Markman & Wachtel, 1988; Soja, Carey & Spelke, 1985; Waxman & Gelman, 1986; Waxman, 1989). A sophisticated intelligent adult, let alone a 2-year-old, would never be able to settle on the meaning of a word by openmindedly considering every possible hypothesis and waiting for evidence to come in that would be sufficient to rule out all but one. Several biases, assumptions, or constraints on word meaning have been suggested as ones used by children to narrow their hypotheses about word meanings: the *whole object*, *taxonomic*, and the *mutual exclusivity* assumptions.

THE TAXONOMIC AND WHOLE OBJECT ASSUMPTIONS

When an adult points to an object and labels it, the novel term could refer to an object category, but it could also refer to a part of the object, or its substance, or color, or weight, among other things. As just mentioned, it is very unlikely that children wait until enough evidence has accumulated to decide among the alternative hypotheses. Instead, one way children initially constrain word meanings is to assume that a novel label is likely to refer to the whole object and not to its parts, substance, or other properties.

Once children decide a term refers to the whole object, they still need to decide how to extend it to other objects. The term could refer to some external relation between two objects. Spatial relations, causal relations, possessor–possessed are some examples of common relations between objects that a term could in principle label. More generally, objects can be related through the variety of ways in which they participate in the same event or theme (e.g., cats eat mice; people read books; birds build nests). Many studies of classification in children demonstrate that children often find thematic relations particularly salient and interesting (see Gelman & Baillargeon, 1983; Markman, 1989; and Markman & Callanan, 1983 for discussions). Having a powerful thematic relation between two objects does not, however, render them the same kind of thing.

If children are attending to thematic relations between objects, how is it that they so readily learn labels for kinds of objects instead? To answer this question, Markman and Hutchinson (1984) proposed that children constrain the possible meanings of words such that they rule out thematic

meanings. That is, children reject thematic relations as a first hypothesis about what a novel label might refer to, despite the fact that they find such relations to be salient and interesting. Markman and Hutchinson conducted a series of studies which compared how children would organize objects when an object was referred to with a novel label versus when it was not. When presented with two objects, such as a dog and cat, and a third object that was thematically related such as dog food, children would often select a dog and dog food as being the same kind of thing. If, however, the dog was called by an unfamiliar label such as *dax* and children told to find another *dax*, they now were more likely to select the cat. These findings have been extended and refined in a number of different studies (Hutchinson, 1984; Soja, Carey & Spelke, 1985; Waxman & Gelman, 1986; Waxman, 1990; Waxman & Kosowski, 1990). Thus, when children believe that they are learning a new *word*, they shift their attention from thematic to taxonomic organization.

MUTUAL EXCLUSIVITY

The whole object assumption leads children to interpret novel terms as labels for whole objects—not for parts or substances of objects or for other properties. But children must of course learn terms that refer to parts, substances, and other properties. The mutual exclusivity assumption, which leads children to expect that each object will have only one label, helps children override the whole object assumption, thereby enabling them to acquire terms other than object labels.

To see how mutual exclusivity overrides the whole object assumption and helps children acquire property terms, suppose a novel term is applied to an object for which a child already has a label. In order to adhere to the principle of mutual exclusivity, the child would have to reject the novel term as a label for the object. The child could simply reject the term as a label for the object without coming up with an alternative meaning. Rejecting one meaning for the term, however, leaves the child with a term that is not yet attached to any referent. This in itself may motivate children to try to find some meaning for the novel term. The mutual exclusivity principle does not speak to how children select among the potential meanings, but they might analyze the object for some interesting part or property and interpret the novel term as applying to it. Studies 2–6 of Markman and Wachtel (1988) demonstrated that 3- and 4-year-old children can use mutual exclusivity to learn terms for parts and for substances. When a novel label was mentioned in the presence of an object with a known label, children rejected the term as a second label for the object and interpreted it instead as a label for a part of the object or its substance.

In addition to overriding the whole object assumption, mutual exclusivity can provide an indirect strategy for acquiring word meanings. Suppose a child observes two objects, one of which already has a known label and one of which does not. If a new label is then mentioned, the child should: (1) on the whole object assumption, look for an object as a first hypothesis about the meaning of the label; (2) on the mutual exclusivity assumption, reject the already labeled object; and therefore, (3) assume the other object is being referred to by the novel label. In this way, the mutual exclusivity assumption enables children to learn the referent of a term without anyone ever explicitly pointing it out. Several recent studies have found that young children can learn object labels by such indirect means (Au & Glusman, 1990; Dockrell & Campbell, 1986; Golinkoff et al., 1985; Hutchinson, 1986; Markman & Wachtel, 1988).

Mutual exclusivity could further contribute to word learning by helping children to narrow overextensions (Barrett, 1978; Clark, 1983, 1987; Merriman & Bowman, 1989). Suppose a child has overextended *dog* to apply to sheep as well as dogs, but then learns the correct name for sheep. The child would then need to stop calling sheep *dog* in order to avoid having two names for the same object. In this way, learning new, correct names for things reduces overextensions.

Clark (1983, 1987) postulates another related principle to help account for semantic acquisition. She argues, following Bolinger (1977), that every word in a dictionary contrasts with every other word and that to acquire words children must assume that word meanings are contrastive. Mutual exclusivity is one kind of contrast, but it is a more specific and stronger assumption: many terms that contrast in meaning are not mutually exclusive. Terms at different levels of a class-inclusion hierarchy, such as *dog* and *animal*, contrast in meaning in Clark's sense, since obviously the meaning of *animal* is different from that of *dog*. Yet, these terms violate mutual exclusivity. However, some of the evidence that Clark (1987) cites for the principle of contrast is, in fact, evidence in support of mutual exclusivity as well. (See Markman, 1989 for a comparison of lexical contrast and mutual exclusivity.)

In addition to the whole object and taxonomic assumptions, then, children constrain word meanings by assuming at first that words are mutually exclusive—that each object will have only one label. Given the nature and function of category terms, they will often tend to be mutually exclusive. A single object cannot be both a cow and a bird or a cow and a dog. Of course there are exceptions: categories can overlap, as in *dog* and *pet*, and they can be included, as in *poodle* and *dog*. Thus, mutual exclusivity is a reasonable, though not infallible, assumption to make. Sometimes, then, children will be led astray by assuming terms to be mutually

exclusive. Adhering to this assumption thus helps explain why children find class inclusion difficult—because it violates mutual exclusivity (Markman, 1987, 1989).

Although there are disadvantages to assuming that object labels are mutually exclusive, the advantages are that by assuming mutual exclusivity, children could avoid redundant hypotheses about the meanings of category terms, narrow overgeneralizations of terms, infer the correct referent of a term without anyone explicitly pointing it out, and override the whole object assumption.

The context in which Merriman and Bowman's monograph appears is one of controversy about whether postulating constraints on word learning is a useful way to conceptualize the problem (Gathercole, 1989; Nelson, 1988), and more specifically whether children use mutual exclusivity to guide their hypotheses about the meanings of novel words. There has been disagreement as to whether or not children possess a mutual exclusivity bias at all. Empirical evidence has been offered to support both sides of the debate. Merriman and Bowman's work makes the important contribution of reviewing this evidence and providing a conceptual analysis of the bias which fits the available evidence. They posit two key aspects of mutual exclusivity: first that it acts as a default option and second that it can be maintained in several different ways.

CONSTRAINTS AS DEFAULT ASSUMPTIONS

Merriman and Bowman treat the mutual exclusivity constraint as a default assumption. There are two components to this argument. One is that in the absence of information to the contrary, children will assume mutual exclusivity but that mutual exclusivity can be overridden. The second is that mutual exclusivity should be treated as a probabilistic bias that guides children's hypotheses rather than as absolute. We agree completely with both of these points and think that they are important enough to amplify. We extend their conceptualization beyond mutual exclusivity to a claim about how constraints in general can guide word learning.

This claim that constraints should be conceptualized as probabilistic biases and default assumptions has been challenged recently by Nelson (1988), who argues that it is inconsistent to hold on the one hand that there are constraints and on the other hand that there are exceptions or violations of the constraints (see also Gathercole, 1989). For Nelson, constraints must be absolute. She argues that any deviation in a child's performance is evidence against a constraint operating. For example, she criticizes the Markman and Hutchinson (1984) evidence that children honor the taxonomic assumption on the grounds that children in these studies were not scoring 100% correct. Nelson's view, then, is that to argue that there may be constraints on word learning dictates a position

that these biases are absolute, admitting of no variance. This is certainly not the position of researchers who have proposed constraints on learning for domains such as conceptual development (Keil, 1979), causal reasoning (Brown, 1990; Gelman, 1990a), counting (Gelman, 1990a), the development of knowledge about objects (Spelke, 1990), and language acquisition (Carey, in press; Markman, 1987, 1989, 1990, in press; Markman & Hutchinson, 1984; Markman & Wachtel, 1988; Newport, 1990; Pinker, 1984; Waxman, 1989) (see also Keil, 1981, 1990, and Gelman, 1990b). Moreover, as we summarize next, the notion of constraints as absolute is not held by ethologists arguing for biological bases of learning (cf. Marler & Terrace, 1984).

Nelson's position might be clarified by the distinction made by Mayr (1974) between *closed* and *open* genetic programs. One way of determining animal behavior is through a closed genetic program, one which does not allow appreciable modification through experience. These genetic programs would be absolute in Nelson's sense. Open programs, in contrast, allow for additional input and modification—they allow for learning. Mayr argues that open programs are more likely in species with longer lifespans and substantial parental care which provide opportunities to learn from experience. Yet Mayr takes care to point out that "an open program is by no means a tabula rasa: certain types of information are more easily incorporated than others" (p. 652), and he cites work on specific hunger and preparedness in learning to make this point (Garcia, McGowan, Ervin, & Koelling, 1968; Rozin & Kalat, 1971). In Mayr's terminology, then, constraints on word learning would clearly not qualify as a *closed* program. If this is the point that Nelson was making then, of course, we would agree. Yet the idea of constraints as probabilistic biases is completely consistent with the way this notion is treated in current ethological theory.

One way in which biases are not absolute is that they may be ordered into a hierarchy such that one bias overrides another. The extraordinary ability of migratory birds exemplifies such a case. Keeton (1974) summarizes some of the most impressive of the documented feats of such birds. A manx shearwater, for example, migrated over 3000 miles in 12½ days to return to its burrow (Matthews, 1953 as cited in Keeton, 1974). In studying homing pigeons, Keeton concludes that when the sun is visible the pigeons use it as a compass. On overcast days, however, the pigeons are still able to find their way home. Thus, the birds have some alternative means of navigation that serve as a back-up system. Keeton reviews the controversy about whether pigeons could be using the earth's magnetic field as one such system. Although this hypothesis was first put forward in 1882 and revived in 1947 there was so much contradictory evidence that it fell into disrepute. The reason for the failures to find that pigeons could

navigate by geomagnetism is that it is not the birds' preferred strategy. Only when the preferred cue for navigation (the position of the sun) is unavailable do pigeons resort to relying on magnetism.

Imprinting provides a good example of a system of substantial plasticity that is nevertheless governed to some extent by innate predispositions. A given species of bird can be sexually imprinted onto a different species or even, in the case of hand-reared birds, onto humans. The birds will later show mating displays towards the foster species. On the other hand, Immelmann (1972) documents that, despite this plasticity, there are preferences for a member of a bird's own species. In a test of whether zebra finches imprint most easily on their own species, male zebra finches were raised by a mixed pair of foster parents, a zebra finch and a Bengalese finch. Although there was equal opportunity for imprinting on either species, the birds nearly always had a sexual preference for their own species. Furthermore, imprinting onto a member of a bird's own species occurs more quickly, is more rigid, and is less likely to be reversed than imprinting onto a different species. Thus the ease of learning and the quality of learning through imprinting is governed in part by the species-specific biases of the animal.

In a recent conference designed to consider issues of constraints on learning in biology, this point that constraints should be thought of as probabilistic biases was made repeatedly (Marler and Terrace, 1984). Here is an example from Gould and Marler (1984) who argue:

Indeed, it is tempting to place a default value interpretation on the associative biases of animals. Although bees, for instance, can learn that a flower is any color from yellow to ultraviolet, they learn the color of purple flowers far more quickly than any other color of flowers (Gould, 1984). At the same time, bees prefer purple silhouettes to all other colors on a spontaneous preference test. It is as though purple is the default parameter—a probabilistic bias which helps guide bees when they experiment with various flowers while searching for food. (p. 65)

And from Gould (1984):

In a very real sense, many cases of selective learning should be thought of as mechanisms by which experience serves to tune an animal's behavior from the default distribution of alternatives to the actual odds in the world around it. (p. 153)

Among ethologists, constraints are postulated as one means of helping the organisms to solve the inductive problems they face. In many cases, these biases do not and could not provide absolute guarantees of correct answers. The environment is too unpredictable for absolute biases to be adaptive. Rather the organism must be capable of learning—of extracting information from the environment. These biases give the organism a good first guess—a head start in solving the problem, compared to if it were sampling randomly from an extraordinarily large number of options.

It is in this way that constraints may be useful for young children trying to figure out what words in their language mean. The constraints that have been postulated, such as the whole object and taxonomic assumptions, in addition to mutual exclusivity, give the young child good first guesses about the meaning of a novel term. They provide powerful means to begin word learning—but not at all the final solutions. Along with mutual exclusivity, these other assumptions should be treated as default assumptions and probabilistic biases.

Take the whole object assumption, for example: without evidence to the contrary, children should interpret a novel term as a label for an object—rather than a part or substance of the object or its color, size, shape, weight, etc. Several different kinds of information could provide evidence to the contrary. If there were no salient object around at the time a novel term was introduced, the absence of a candidate object could override the whole object assumption. For example, Soja et al. (1985) found that when presented with a blob of stuff rather than a discrete object, children will interpret a novel label as a substance term. As Markman and Wachtel (1988) demonstrated, the mutual exclusivity assumption can be another source of information in conflict with the whole object assumption. By rejecting a novel term as a second label for an object, children will then search for a part, or substance, or other attribute of the object to label. Thus children will violate the whole object assumption in order to preserve mutual exclusivity. As children learn more about their language, grammatical form class can serve as a further means of overriding the whole object assumption. If for example the novel word is clearly recognizable as a verb, that would cause children to override the whole object assumption. In sum, the whole object constraint serves as a first hypothesis that can be overridden in a variety of different ways ranging from lack of environmental support (e.g., when there are no salient objects around) to its coming into conflict with other word learning constraints (e.g., mutual exclusivity) to its conflict with other aspects of the linguistic system (e.g., grammatical form class).

Similarly, while children's first hypothesis should be that terms are mutually exclusive, this bias can be overridden by evidence to the contrary. The ability to overcome mutual exclusivity is important because children must eventually learn words that are not mutually exclusive (e.g., superordinate and basic level terms). One illustration of the kind of information that is used by children comes from Gelman, Wilcox, and Clark (1989), who taught 3- through 5-year-olds second labels for objects but varied whether the second label was a simple label or a compound noun (such as *oak-tree* or *taxi-car*). They reasoned that the compound form might provide a linguistic clue that the objects were to be labeled at more than one hierarchical level. Under these circumstances

children made quite a few errors in learning the second label, and the majority of the errors they made were to treat the labels as mutually exclusive subsets. However, hearing compound nouns as second labels helped children learn two labels for the same object especially at the subordinate level. By explicitly representing both levels of the hierarchy, compound nouns provide one way of helping children override mutual exclusivity.

One implication of viewing constraints as default assumptions is that violations of a constraint found in a child's lexicon are not necessarily evidence against the existence of the constraint. Yet such counterexamples constitute much of what has been taken as evidence against constraints (Banigan & Mervis, 1988; Gathercole, 1987, 1989; Merriman, 1987; Mervis, 1987, 1989; Nelson, 1988). Instead of treating such violations simply as negative evidence, we could look to such violations as information about how children go about overriding the constraint when needed. Merriman and Bowman approach the literature on mutual exclusivity from this perspective. To keep the default assumption interpretation as a testable hypothesis, then, one must be able to determine whether violations of mutual exclusivity are to be taken as evidence against the existence of the constraint, or whether they should be taken as instances in which the default assumption was overridden. To answer this, Merriman and Bowman elaborate on what is required to document a child's failure to use mutual exclusivity. In addition to arguing that mutual exclusivity is a default assumption, they argue that there is flexibility in how it can be maintained. Even when mutual exclusivity is preserved, children are not restricted to one set response, but, rather, are able to make use of different aspects of the situation to maintain mutual exclusivity.

WAYS MUTUAL EXCLUSIVITY CAN BE MANIFESTED

Merriman and Bowman outline four ways in which children can act in accord with mutual exclusivity: *disambiguation*, *correction*, *rejection*, and *restriction*. If a new term is used in a context in which it could either refer to an object with a known label or one whose label is not yet known, children should avoid interpreting the term as a second label for the known object and interpret it instead as referring to the object they cannot name. Merriman and Bowman call this the disambiguation effect. Alternatively, when presented with a second label for an object, a child could correct the old label, replacing it with the new one. Another option would be to simply reject the second label, either by explicitly denying that the term is appropriate (e.g., "No, that's not a . . .") or by just ignoring the second label. Finally, in order to preserve mutual exclusivity, children should always avoid generalizing a label to already named items, that is, show the restriction effect. Merriman and Bowman note that which of the

options for maintaining mutual exclusivity is used depends on the situation. If the reference of the second label is ambiguous, the child is likely to disambiguate; that is, to map the label onto an object (or part, or property) without a known name. If the child is uncertain about the old name, he or she may correct it, replacing it with the new one. Or the child might simply reject the new name. Merriman and Bowman conclude that there are only two strict implications of acting in accord with mutual exclusivity: first that children will either disambiguate, correct, or reject, and second that they will always restrict.

PROBLEMS WITH INFERRING THE PROCESS FROM THE LEXICON

This view of mutual exclusivity as a default option that can be manifested in a variety of ways implies that violations of a constraint found in a child's lexicon do not necessarily invalidate the constraint. As mentioned earlier, such violations have been interpreted as refuting mutual exclusivity. However, the existence of violations is not sufficient to show that children lack the bias. How the interpretation was arrived at is what is at issue, not only what was acquired. For example, by postulating the whole object assumption, one is not committed to a position that says children are incapable of learning property terms and that if one finds an adjective in a child's vocabulary the constraint is disproved. Rather, the argument is that object labels will typically constitute children's initial hypotheses upon hearing a novel word, and in order to learn property terms children must override that initial bias. Similarly, to claim that children are biased to treat object labels as mutually exclusive is not to claim that they can never learn more than one label for the same object. The test of the hypothesis requires examining the order of hypotheses children consider to see whether they resist violating mutual exclusivity on first hearing a novel word. If a child's initial hypothesis reveals an attempt to preserve mutual exclusivity then that would argue in favor of mutual exclusivity as a constraint on word learning even if the child is ultimately successful at overriding the constraint.

MERRIMAN AND BOWMAN'S ANALYSIS OF THE EXISTING EMPIRICAL WORK

In sum, Merriman and Bowman have carefully analyzed what is required to disprove the existence of a mutual exclusivity bias which is a default option and potentially manifest in different ways. They review the literature on mutual exclusivity and lexical contrast from the perspective gained by their analysis. Their thoughtful review is detailed and thorough, and covers an extensive literature. We will not attempt to summarize it here except to make a few very general points.

Critique of Evidence Against Mutual Exclusivity

There are two broad kinds of criticisms Merriman and Bowman make about the putative counterevidence to mutual exclusivity. The first has been foreshadowed by our earlier discussion, namely that the counterevidence consists of examples of violations of mutual exclusivity found in children's lexicons. For example, diary studies of children's language acquisition usually report a few cases of words overlapping in reference: Merriman (1987) notes that there are nine violations of mutual exclusivity in Leopold's (1939, 1949) corpus, as well as several instances of violations in less extensive diary records. In an observational study, Mervis (1987) found that her subject continued to use overextended names for some objects after they had begun to use the correct terms for them. Findings like these have been taken to refute mutual exclusivity, but, as noted earlier, they reflect the conclusion of some process of word acquisition and not the process itself. We cannot judge what hypotheses children may have begun with when the only record is the end product, the words that they have learned.

Merriman and Bowman's second general criticism of existing studies is that because there are a number of ways in which children can act in accord with mutual exclusivity, it is important to check all the possibilities before concluding that they are not doing so in a given situation. Existing studies have failed to do this. For example, Mervis (1987) concluded that young children do not have the bias based on their failure to correct this mistaken "child-basic" name once they accepted the adult word. Yet she notes that: "there were many instances in which a child vigorously rejected the mother's attempt to introduce . . . a new name for an object the child already included in a child-basic category labeled by a different name" (p. 223). Although children did not correct their label, they did show the rejection effect in initially resisting a second label, thus giving evidence of mutual exclusivity on Merriman and Bowman's analysis.

Critique of Evidence Supporting Mutual Exclusivity

Merriman and Bowman also raise an important criticism of some of the evidence in support of mutual exclusivity, in particular, they suggest an alternative explanation for the findings of studies such as those by Golinkoff et al. (1985), Hutchinson (1986), and Markman and Wachtel's (1988) first study. These studies focused on what Merriman and Bowman call the disambiguation effect. As mentioned earlier, in the presence of two objects, one of which has a known name and one of which does not, children will interpret a novel label as a label for the novel object. This has been taken to support children's use of mutual exclusivity in that children

are presumed to be rejecting the label as a second label for the known object and thus interpreting it as a label for the as-yet-unnamed object. Merriman and Bowman quite rightly point out that these results could be obtained without recourse to mutual exclusivity if children had a bias to fill lexical gaps. The lexical gap hypothesis states that in the presence of an object that as yet has no known label, children are motivated to discover its name (cf. Clark, 1983, 1987). The assumption that things have names is related to the development of the "nominal insight" that McShane (1979) and others have postulated. Many children early in the course of language acquisition actively request labels for things. Thus, in the tests of mutual exclusivity just described children could map the novel word to the novel object because they have a novel object that they want to find a name for. If children had no reluctance to have second labels for things but simply were motivated to find first labels for things, then the results of these studies would be the same. Thus the lexical gap hypothesis might explain some of the evidence that has been interpreted as evidence in favor of mutual exclusivity. Whether children are in fact using mutual exclusivity in these cases is an open question.

Merriman and Bowman extend the lexical gap hypothesis to cover the results of Markman and Wachtel's (1988) studies which demonstrate how mutual exclusivity can enable children to learn terms for parts and substances. We find this part of their argument less compelling. Recall that Markman and Wachtel found that children were better able to learn a part term or a substance term when it was attributed to objects that already had a known label. On the whole object assumption, children who hear a novel label should first interpret it as a label for the object as a whole, not for one of its parts, its substance, or any other property. But when the term is attributed to an object with a known label, on the basis of mutual exclusivity children should reject the term as a second label, thus overriding the whole object assumption, and should be motivated to analyze the object for some other property to label. A concrete example is that children who heard "this is pewter" attributed to a cup (an object whose label was known) were more likely to think that *pewter* referred to the substance and not the object than were children who heard *pewter* attributed to tongs (an object whose label was unknown). Here is how the lexical gap hypothesis can explain these results. The argument is that children who hear a cup called *pewter* have no reluctance to treat *pewter* as a second label for cup. But they prefer to fill lexical gaps over having redundant labels. Children seek a label for the substance of the cup, because they don't yet know one, so they prefer to treat *pewter* as a label for the substance rather than as a second label for cup.

The difference between the two accounts is subtle. On the mutual exclusivity account it is the rejection of the whole object as a possible

referent of the term that motivates children to try to find something else as the referent of the term. On the lexical gap hypothesis, the child has noticed parts and properties without labels and waits for an opportunity to label them. Although as yet there is no clear evidence to distinguish between these two positions, a number of anecdotal reports favor the mutual exclusivity hypothesis. Several investigators (Clark, 1987; Macnamara, 1982; Mervis, 1987) have reported that young children on first hearing a novel label attributed to an object with a known label, simply reject the second label. For example a child who is told "see the poodle" will insist that it is a dog and not a poodle. This of course is readily explained by assuming children reject second labels for objects. On the lexical gap hypothesis, however, children are assumed to be motivated to find labels for as yet unlabeled parts and properties. Why then should they simply reject the term instead of interpreting it as a label for paw, or tail, or curls, or fur? In our judgment mutual exclusivity does a better job of explaining why children are better able to learn property terms when they are attributed to objects with known labels.

In sum, Merriman and Bowman's analysis establishes several guidelines for empirically testing the constraints on word learning hypothesis. Before concluding that a child has violated mutual exclusivity, one must check for all of mutual exclusivity's potential manifestations. Mutual exclusivity is not an unchangeable, preprogrammed response, but rather a bias which can influence children's hypotheses about word meanings differently in different situations. Furthermore, when children do violate mutual exclusivity, it is important to explore the factors that may have led them to do so. Evidence for mutual exclusivity ought to be seen in the process of word learning, not in the static lexicon. Finally, Merriman and Bowman present a compelling alternative explanation for some of the evidence in favor of mutual exclusivity, namely that a propensity to fill lexical gaps will also predict that children will map novel labels onto objects with no known name rather than to objects whose label is already known. With these conclusions drawn from Merriman and Bowman's review, we now turn to their experimental studies of mutual exclusivity in young children.

MERRIMAN AND BOWMAN'S STUDIES

Following their account of mutual exclusivity, Merriman and Bowman explored empirically the roles the bias plays in children's language learning. In keeping with their model, they looked for several potential manifestations of mutual exclusivity, designing experimental situations which ought to lead children to take different courses of action to maintain mutual exclusivity. Merriman and Bowman's research included 2-year-olds as well as preschoolers and older children. We will begin by

discussing their findings for preschoolers, and then discuss their conclusion about 2-year-olds, which we find problematic.

Findings with Children 2½ and Older

Tests of disambiguation. One way that children may show mutual exclusivity is to use the bias to eliminate objects with known labels as potential referents for a new label. Merriman and Bowman report several versions of a disambiguation task, in which children are presented with two objects, one novel and one familiar, and asked which is the referent of a novel word. If children follow mutual exclusivity, they ought to rule out the familiar object as a potential referent and choose the novel object. As mentioned earlier, this kind of task has been used by a number of other researchers to demonstrate mutual exclusivity in children (Au & Glusman, 1990; Dockrell & Campbell, 1986; Golinkoff et al., 1985; Hutchinson, 1986; Markman & Wachtel, 1988).

Merriman and Bowman found that preschoolers showed the disambiguation effect, thus confirming the findings of previous research. However, given Merriman and Bowman's earlier criticism of this kind of disambiguation task, it is surprising that they chose to use it as a test of mutual exclusivity. Recall that Merriman and Bowman correctly point out that the findings of standard disambiguation tasks are confounded with the possible effects of a bias to fill lexical gaps. Thus, while Merriman and Bowman's results are in keeping with earlier findings, whether they are evidence for mutual exclusivity remains uncertain.

Tests of correction and rejection. When a second label is given unambiguously, children could honor mutual exclusivity by two means: they could correct the previously held name or reject the second label. While some studies have shown that children sometimes correct first labels (Banigan & Mervis, 1988; Chapman, Leonard & Mervis, 1986; Merriman, 1986) and others report anecdotal evidence that children often reject second labels (Clark, 1987; Macnamara, 1982; Mervis, 1987), none have systematically varied the conditions that lead children to take one option rather than the other. Merriman and Bowman propose that children's choices may be influenced by how certain they are of the object's first label. In their studies of correction and rejection, they manipulated the extent to which children were certain of an object's name, predicting that when children were certain of the object's first label they would reject the second label, but that when they were uncertain, they would replace the old label with the new one.

Merriman and Bowman designed a task in which children were shown sets of six pictures of objects and told a novel name for one of them. Each set of pictures consisted of two typical exemplars of two kinds from the same semantic field (e.g., two trucks and two cars), and one hybrid which

had features of both kinds (e.g., car-truck), and one unrelated item (e.g., a hand). The sets were designed so that children would be certain of the names for the typical exemplars, but less certain of the correct name for the hybrid items. Children were told an invented name (e.g., *jegger*) for one of the pictures. Half the children were told that the hybrid was a *jegger*, the other half were told that one of the typical objects was a *jegger*. Then, the rest of the set of six was brought out and the child was asked to show the experimenter which ones were *cars* and which were *trucks*. When the children are certain of the old name (i.e., when the labeled object is a typical exemplar), they should reject the new label, but when they are unsure (i.e., for the hybrid object) they should correct the label they had originally held for that object and accept the new label for it. Thus, in the test phase, children who were told that the hybrid was a *jegger* should be less likely to include it as a *car* or a *truck* than when they had heard no object labeled. The children who were told that the typical car was a *jegger* should not show this difference. The predicted pattern of results was found for the preschool-aged children. Thus, for these children, certainty about an item's label determined which strategy they selected to preserve mutual exclusivity. These results show not only that children are motivated to honor mutual exclusivity, but also that they do so differently depending on contextual factors.

Tests of restriction. The results from the tests of correction and rejection provided information on whether children showed the restriction effect as well. To preserve mutual exclusivity children should not allow object labels to overlap. By examining children's responses when asked for the *cars* and *trucks* in the tests of correction, Merriman and Bowman were able to assess the extent to which the children violated the restriction rule. If children chose overlapping sets as the referents for the two labels (i.e., chose some items as both *cars* and *trucks*) it was concluded that they had violated mutual exclusivity. As predicted, preschoolers tended to avoid choosing overlapping sets for the two names.

Merriman and Bowman's results demonstrate the mutual exclusivity guides children's word learning in a number of different ways. When there is another likely object present, children disambiguate (possibly via mutual exclusivity, but possibly through their motivation to fill lexical gaps). When a second label is given unambiguously to an already named object, mutual exclusivity leads children sometimes to correct the old label, and sometimes to reject the new one. Moreover, it motivates preschoolers to keep word extensions from overlapping, thus showing the restriction effect.

Findings with 2-Year-Olds

As mentioned earlier, Merriman and Bowman looked at mutual exclu-

sivity in children of varying ages. In several of their studies, they included not only preschool-aged children, but also children as young as 2 years. In all of the studies in which 2-year-olds were included, they failed to show the effects predicted by having mutual exclusivity. That is, they appeared to allow objects to be referred to by more than one label. From this, Merriman and Bowman conclude that children below the age of 2½ do not have a mutual exclusivity bias. We disagree with this conclusion. The tasks they used were unsuitable for very young children in a number of ways, for example in failing to equate the amount of time children were allowed to play with novel versus familiar objects, and in having lengthy test sessions requiring children to learn a great many novel words at once. These procedural problems make it impossible to conclude that lack of mutual exclusivity was responsible for the 2-year-olds' failure. Moreover, there is evidence that mutual exclusivity is in fact used by 2-year-olds. Before detailing these specific objections, we first consider how relevant age of onset is to questions about the origins and usefulness of mutual exclusivity. Even if Merriman and Bowman were correct that mutual exclusivity appears only after age 2½ this would not invalidate claims as to its usefulness nor would it bear on the question of whether the constraint is innate.

Issues surrounding age of onset. The age at which children can first make use of mutual exclusivity has important implications for accounts of early language acquisition, however, it is important to be clear about what these implications are. Some researchers see age of onset as evidence indicating whether a behavior is innate or learned. They hold that what appears early is likely to be innate, while what appears late is likely to be learned. If mutual exclusivity or other constraints do not appear until children are 2 or 3, they contend, then these constraints must be learned and not innately given. Nelson (1988) favors the "learned" side of the debate, and points to evidence for late onset as supporting this view. Similarly, MacWhinney (1989), in his commentary to the monograph, holds that "mutual exclusivity is not a fundamental constraint, but a set of learned assumptions" (p. 129), and he hails Merriman and Bowman's conclusions about age of onset as "turning the argument for an innate mutual exclusivity bias on its head" (p. 126).

However, age of onset alone cannot resolve the learned versus innate debate. Many innate abilities are late emerging, and many learned abilities are acquired early. No one would seriously question the biological basis of puberty even though it does not occur early in life. It is too soon yet to say whether mutual exclusivity and other constraints owe a larger debt to innate factors or learning. The age of onset findings, however, do not disprove the innateness hypothesis—they are uninformative on this issue.

Moreover, conceptualizing the issues surrounding the origins of con-

straints as if they boiled down to innate versus learned oversimplifies the question (Bateson, 1984; Mayr, 1974). At this early stage of our knowledge of word-learning constraints, an innate-learned dichotomy obscures many interesting developmental possibilities. For example, a given constraint could be the by-product of several independent processes, each with its own developmental history, which converge on language learning. The question of the origins of the mutual exclusivity bias is an important one, but it is as yet unanswered.

Shettleworth (1972, 1984) states another way in which the innate versus learned dichotomy is an oversimplification. She argues that learning itself should be examined from an evolutionary perspective and that investigators should view learning as an adaptation by considering issues such as:

Why learn in a specific case as opposed to relying on other solutions, including that of not making the adjustment at all?

What to learn? For example, when an individual must learn to recognize individual conspecifics or the approach of a predator, what cues does he use? Are they the optimal ones in the sense of being the most reliable predictors in the situation?

When to learn? Does learning begin immediately on first exposure to a situation? Does general learning about the environment occur during periods of 'sampling' or play and get put to use when it is needed? . . .

How to learn? Trial and error, imprinting, observation, association: is a particular process the only one that can solve the problem, or if several might serve the purpose, what determines which one a given species uses?

How fast to learn and how long to remember? (Shettleworth, 1984, p. 429)

Although the discussions of the innateness of word-learning constraints has been oversimplified and age of onset cannot address the questions of the origins of the constraint, age of onset *can* tell us what tools children have available for use at different stages in their development of language, and this can inform our accounts of the mechanisms of language acquisition. Age of onset is important for knowing what children at the earliest stages of learning language have to help them. There is evidence that children use the taxonomic and whole object constraints quite early, at 18 months or younger (Backscheider & Markman, in preparation; Huttenlocher & Smiley, 1987). These constraints would be very useful to children who are just breaking into language because they limit the number of hypotheses children consider when figuring out the meaning of a new word. As a framework for organizing the beginning lexicon, mutual exclusivity could also be useful in this capacity and might well be in place from the earliest stages of language acquisition.

On the other hand, mutual exclusivity may not be needed in the earliest phase of language acquisition when children have extremely limited vo-

cabularies. If mutual exclusivity is thought of as a method of systematizing knowledge (Markman & Wachtel, 1988), there may be reason to expect it to appear after at least some language has been learned. In order for children to be motivated to systematize knowledge in a domain, they may first need to have acquired enough knowledge to warrant systematization (Karmiloff-Smith, 1979; Karmiloff-Smith & Inhelder, 1975). Mutual exclusivity would then be needed by children who have already begun to acquire language. As already mentioned, not only would it help children to limit their redundant hypotheses about the meaning of a new word and to narrow overextensions, it could provide an indirect strategy for acquiring object labels, and by overriding the whole object bias, it could motivate the child to learn part and property terms.

Although mutual exclusivity would be useful to very young language learners, Merriman and Bowman propose that children lack mutual exclusivity before the age of 2½. They suggest that other factors, such as parental labeling strategies and corrective feedback, could substitute for mutual exclusivity. We are skeptical that the kinds of factors Merriman and Bowman discuss are sufficient to account for language learning up to the age of 2½. Whether, in principle, parental strategies could accomplish what mutual exclusivity can, and, if so, whether parents are consistent in their use of such strategies, are questions that have by no means been answered satisfactorily. A straightforward way to approach the question is to see whether children younger than 2½ show any signs of having the bias. In their empirical work, Merriman and Bowman addressed this question by evaluating the performance of children of varying ages on their tests of mutual exclusivity. However, these tasks have problems that make them inconclusive on the issue of age of onset.

Problems for the tests of disambiguation. Although children of 2½ or 3 years chose the novel object as the referent of the novel label in Merriman and Bowman's disambiguation tasks, 2-year-olds did not. Merriman and Bowman thus conclude that 2-year-olds have no bias to treat terms as mutually exclusive. Recall, however, that Merriman and Bowman have argued that children may map a novel label to a novel object on disambiguation tasks not because they have a mutual exclusivity bias, but because they have a bias to fill lexical gaps. Given this, children's failure to show the disambiguation effect would lead to the unlikely conclusion that they lack a bias to fill lexical gaps. This is problematic because children seem quite motivated to learn the names for the objects around them from an early age. At around 18 months, during the "naming explosion," children become intensely interested in learning what as-yet-unnamed things are called (Bloom, 1973; Gopnik & Meltzoff, 1986, 1987; Nelson, 1973, 1988). That Merriman and Bowman's studies did not pick up on this

well-known tendency suggests that their procedures were insensitive to the younger children's knowledge. They may have failed to detect the 2-year-olds' competence because of problems in the design of their studies.

One problem arises from how Merriman and Bowman attempted to control for novelty. Children might choose the novel object in disambiguation tasks simply because they like it better. Different researchers have dealt with this problem in different ways. In order to decrease novelty, Golinkoff et al. (1985) let children play with all of the toys to be used in the task ahead of time. Moreover, in their disambiguation task, children had two different kinds of choices: when first given the novel label, they chose from two familiar objects and one novel one, then they heard the label again, this time with a different novel object as one of the three distractors. If novelty governed the children's choices, they should have chosen the completely novel object on this second trial. They did not do this, but instead were consistent with their earlier choice. Hutchinson (1986) made sure that her subjects did not prefer the novel object before she tested them. Before the disambiguation task she presented children with all of the novel-object/familiar-object pairs twice, asking them to choose the one they liked best. If a child chose the novel object both times, that item was eliminated from the analysis of that child's performance. Very few eliminations were required. Markman and Wachtel (1988) ran a control group in which children were given the same choice as in the disambiguation task, but without a label. In this control condition, children were at chance in choosing between the two objects.

Merriman and Bowman used a different solution to the problem of novelty. Before their tests of disambiguation, they gave children the opportunity to play with the novel objects but not the familiar ones, reasoning that this would eliminate any inherent attractiveness based on novelty. Thus, at the time of testing, children were choosing between an object which they had been playing with for several minutes and one which they had not yet seen. Given well-known habituation effects, the time spent with the novel objects should have made them less interesting to young children. For the 2-year-olds, this exposure seems to have overcompensated for novelty. For example, in Study I, the 2-year-olds chose the object they had been playing with only 29% of the time, which was significantly less than chance. Since this object was always the novel object, this means that they chose the familiar object about 70% of the time. In this case, then, the pre-exposure did not equate the salience of unfamiliar and familiar items, but instead created another imbalance.

Another factor that may have impeded 2-year-olds' performance was the sheer number of trials. In Study I, each child received 16 trials. On each trial a different novel object and novel name were used. Thus, the

children in this study were introduced to 16 new words (more than one a minute) and even more novel objects. It seems likely that this would overwhelm the 2-year-olds. Other tests of disambiguation have not presented children with this density of new information. Markman and Wachtel (1988) gave their 3- to 4-year old subject only six novel labels, Golinkoff et al. (1985) gave their 2½-year-olds six, and Hutchinson (1986) gave her 2- and 3-year olds five. Moreover, in other studies which asked many questions of children at least a few of these were about familiar objects (Golinkoff et al. 1985; Markman & Wachtel, 1988). This allowed the children to feel capable and not lose track of the task. In addition, these trials provided evidence that the children were attending to the task. In Study I, Merriman and Bowman asked only about novel labels, and therefore, there is no way of knowing if the youngest children in this study lacked mutual exclusivity or lacked the stamina to sustain attention to the task.

The findings of the tests of disambiguation in Studies II and IV support the conclusion that the processing demands in Study I overwhelmed the youngest children's true competence. In Study II, Merriman and Bowman decreased the number of novel objects and novel labels to six, and asked about familiar objects on another six trials. This study revealed a knowledge of mutual exclusivity that was missed in Study I: in Study I, 3-year-olds were at chance in their choices on disambiguation with exposure trials; in Study II, even the 2½-year-olds were well above chance on these trials. In Study IV, children were given only three trials, and here as well, the 2½-year-olds were correct in their choices. Unfortunately, no 2-year-olds were run in this study.

From these studies we cannot conclude that children under the age of 2½ do not disambiguate. There are a number of likely explanations for the young children's poor performance which have nothing to do with a lack of mutual exclusivity. Furthermore, Hutchinson (1986) found that 2-year-olds do show the disambiguation effect. Her youngest subjects, late 1- to early 2-year-olds, were significantly above chance in choosing the novel object, and, as discussed earlier, she adequately controlled for possible effects of novelty.

Problems for the tests of correction and rejection. In Merriman and Bowman's tests of correction and rejection, they found the predicted pattern of results for 2½-year-olds and 3-year-olds: when a second label was given to a typical member of an object category, children seemed to reject the second label; when a second label was given to a hybrid (atypical) object, children were less likely to include the hybrid in *cars* or *trucks*, suggesting that they had corrected their original label for the item by replacing it with the new label. The 2-year-olds did not show this pattern. However, as Merriman and Bowman point out, the logic of this

study requires that children do in fact learn the new name. Unless children learn the name *jegger* for the labeled object, there is no reason for them to decide that that object should not be called a *car* or a *truck*. This precondition was fulfilled for the older children, but not for the younger children. In Study II the 2½-year-olds correctly comprehended 93% of the hybrid items when questioned at the end of each trial. Two-year-olds, on the other hand, comprehended only 40% of the hybrid items. The same pattern of differences was found in Study III. Thus it is likely that the reason labeling did not affect the youngest children's performance is that they did not learn the new labels very well.

The results of the comprehension test show an interesting asymmetry which may reflect a mutual exclusivity bias for the 2-year-olds as well as the older children. Children learned new names for the atypical "hybrid" items more often than for the typical items. That is, if they were shown a typical car, and told it was a *jegger* they appeared less likely to learn the label than if they were taught a label for an odd looking, atypical car-truck. The 2-year-olds learned the new name for 40% of the hybrids but only 27% of the typicals, and the 2½-year-olds learned the new name for 93% of the hybrids but only 50% of the typicals. Thus the extent to which children were certain of an object's name seems to have affected how well they could learn a new name for it. In their review of the literature, Merriman and Bowman point out that it is important to look for all of mutual exclusivity's possible manifestations before concluding that the bias is missing. This is a case in point: while the youngest children did not correct a previously held name, they seem to have shown the rejection effect, which is predicted by Merriman and Bowman's account of mutual exclusivity.

Problems for the tests of restriction. In the tests of restriction, Merriman and Bowman found that although older children chose nonoverlapping sets of items when asked to pick out the *cars* and the *trucks*, 2-year-olds did not. However, here, as in the other studies, the youngest subjects may have found the task confusing. Recall that for the tests of restriction, children were shown sets of six pictures: two pictures each of typical exemplars from two related categories (e.g., cars and trucks), one hybrid (car-truck), and one unrelated distractor. The sets featured: cookies and crackers, cars and trucks, spoons and forks, and shoes and socks. In each set, the two different kinds of items were quite similar, and it is not clear that the youngest children could always tell which was which. While Merriman and Bowman gave the 2-year-olds a test of recognition of the types of items before the task, all that was required for "passing" was for the child to pick an item from the correct semantic field. Thus, a child could choose a cookie when asked for *cracker*, and pass the recognition

test. The 2-year-olds may have been confused about which were cookies and which were crackers. To the extent that children were unsure about the identity of the items, they may have had trouble remembering which ones they had chosen already, and thus chosen overlapping sets out of confusion. Further evidence that the test of restriction was confusing to the young children is that 2½-year-olds did not show the restriction effect, even though they could disambiguate and showed the immediate correction effect.

Thus, the studies presented by Merriman and Bowman are inconclusive as to whether children below the age of 2½ have a mutual exclusivity bias. The evidence Merriman and Bowman rely on is poor performance by 2-year-olds on tasks that are not fair tests for children so young. In addition, although Merriman and Bowman begin by pointing out how important it is to consider all the possible ways mutual exclusivity could be realized before concluding a child lacks the bias, their empirical work falls short of this.

Merriman and Bowman's case is further weakened by suggestions from other research as well as from their own data that children younger than 2½ may well have mutual exclusivity. As noted earlier, Hutchinson (1986) found that 2-year-olds did disambiguate, and in the comprehension phase of Merriman and Bowman's correction test, 2-year-olds appear to reject a second label for an object with a well-known label and learn, at least to some extent, a label for an object whose original label was in doubt. In addition, some work by Mervis and her colleagues (Banigan & Mervis, 1988; Chapman, Leonard, & Mervis, 1986) lends support to the notion of mutual exclusivity in very young children. Although Banigan and Mervis drew the opposite conclusion from their findings, there is an indication in their data that young children have a mutual exclusivity bias. They found that children aged 2 and younger resisted learning a new label for a previously named object. Simply telling the child the second label was not effective in teaching it to children, though we suspect that it would have been enough to teach children the name of a novel object (for evidence on this point, see Liittschwager & Markman, 1991). Rather, the second label had to be accompanied by a verbal description and a demonstration of a salient property of the object in order for the children to accept it. In addition, Mervis (1987), in a longitudinal study of children who were two and younger, found that children resisted learning new (correct, adult) labels for items for which they already had names. Thus, the issue of when children first have mutual exclusivity is unresolved. While Merriman and Bowman do not find evidence of the bias in children under 2½, there are suggestions in the existing literature that very young children do indeed have mutual exclusivity.

CONCLUSION

Merriman and Bowman's monograph is an important contribution to the work on constraints on language learning. In his commentary to the monograph, MacWhinney (1989) suggests that Merriman and Bowman's work has put the notion of a mutual exclusivity constraint to rest. We disagree. Both Merriman and Bowman's theoretical analysis and their empirical work support the idea of constraints on language learning. Chief among their contributions is their elaboration of the idea of mutual exclusivity as a default assumption—a probabilistic bias which can be overridden. This analysis clears up the confusion which has resulted from the misconstrual of constraints as absolute (Nelson, 1988). Treating constraints as rigid, preprogrammed responses, admitting of no variation has led researchers to interpret any violation as evidence against the constraint. The default assumption account, on the other hand, avoids this interpretation. Another contribution of Merriman and Bowman's analysis is that it points out the many ways that the mutual exclusivity constraint can be manifested. This clarifies apparent inconsistencies in the literature in that putative counterexamples often occur when an investigator tests for only one consequence of mutual exclusivity and ignores its other possible implications. Merriman and Bowman's empirical studies document the variety of ways in which mutual exclusivity influences children learning language: possibly leading them to disambiguate, and, given the right circumstances, to correct old names, to restrict overextended names when given a second label, or simply to reject a second label entirely. Although Merriman and Bowman's procedures were too difficult for 2-year-olds thus leaving the question of age of onset unanswered, their findings about the many ways in which mutual exclusivity guides word learning in children from 2½ on confirm and extend the evidence that mutual exclusivity is important for young language learners.

REFERENCES

- Au, T. K., & Glusman, M. (1990). The principle of mutual exclusivity in word learning: To honor or not to honor? *Child Development*, 61, 1474–1490.
- Backscheider, A., & Markman, E. M. (in preparation). *Young children's use of the taxonomic assumption to constrain word meanings*.
- Banigan, R. L., & Mervis, C. B. (1988). Role of adult input in young children's category evolution: II. An experimental study. *Child Language*, 15, 493–505.
- Barrett, M. (1978). Lexical development and overextension in child language. *Journal of Child Language*, 5, 205–219.
- Bateson, P. (1984). Genes, evolution, and learning. In P. Marler & H. S. Terrace (Eds.), *The biology of learning* (pp. 75–88). New York: Springer-Verlag.
- Bloom, L. M. (1973). *One word at a time: The use of single word utterances before syntax*. The Hague: Mouton.
- Bolinger, D. (1977). *Meaning and form*. London: Longman.

- Brown, A. (1990). Domain-specific principles affect learning and transfer in children. *Cognitive Science*, *14*, 107–133.
- Carey, S. (in press). Procedures toddlers use to constrain word meanings—speculations and a little data. In E. Dromi (Ed.), S. Strauss (Series Ed.), *Language and Cognition: The Fifth Annual Tel Aviv Workshop in Human Development*. Norwood, NJ: Ablex.
- Chapman, K. L., Leonard, L. B., & Mervis, C. B. (1986). The effect of feedback on young children's inappropriate word usage. *Journal of Child Language*, *13*, 101–117.
- Clark, E. V. (1983). Meanings and concepts. In J. H. Flavell & E. M. Markman (Eds.), P. H. Mussen (General Ed.), *Handbook of child psychology (Vol. 3): Cognitive development* (pp. 787–840). New York: John Wiley & Sons.
- Clark, E. V. (1987). The principle of contrast: A constraint on acquisition. In B. MacWhinney (Ed.), *Mechanisms of language acquisition: The 20th annual Carnegie symposium on cognition* (pp. 1–33). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Clark, E. V. (1988). On the logic of contrast. *Journal of Child Language*, *15*, 317–335.
- Dockrell, J., & Campbell, R. (1986). Lexical acquisition strategies in the preschool child. In S. Kuczaj & M. Barrett (Eds.), *The development of word meaning* (pp. 121–154). Berlin: Springer-Verlag.
- Garcia, J., McGowan, B. K., Ervin, F. R., & Koelling, R. A. (1968). Cues: Their relative effectiveness as a function of the reinforcer. *Science*, *160*, 794–795.
- Gathercole, V. (1987). The contrastive hypothesis for the acquisition of word meanings: A reconsideration of the theory. *Journal of Child Language*, *14*, 493–531.
- Gathercole, V. (1989). Contrast: A semantic constraint? *Journal of Child Language*, *16*, 685–702.
- Gelman, R. (1990 (a)). First principles organize attention to and learning about relevant data. *Cognitive Science*, *14*, 76–106.
- Gelman, R. (1990 (b)). Structural constraints on cognitive development. *Cognitive Science*, *14*, 3–9.
- Gelman, R., & Baillargeon, R. (1983). A review of some Piagetian concepts. In J. H. Flavell & E. M. Markman (Eds.), P. H. Mussen (General Ed.), *Handbook of child psychology (Vol. 3): Cognitive development* (pp. 167–230). New York, NY: Wiley.
- Gelman, S. A., Wilcox, S. A., & Clark, E. V. (1989). Conceptual and lexical hierarchies in young children. *Cognitive Development*, *4*, 309–326.
- Golinkoff, R. M., Hirsh-Pasek, K., Baduini, C., & Lavalley, A. (1985, October). *What's in a word? The young child's predisposition to use lexical contrast*. Paper presented at the Boston Child Language Conference, Boston, MA.
- Goodman, N. (1955). *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.
- Gopnik, A., & Meltzoff, A. N. (1986). Words, plans, things and locations: Interactions between semantic and cognitive development in one-word stage. In S. Kuczaj & M. Barrett (Eds.), *The development of word meaning* (pp. 199–223). New York: Springer-Verlag.
- Gopnik, A., & Meltzoff, A. N. (1987). Early semantic developments and their relationship to object permanence, means–ends understanding, and categorization. In K. Nelson & A. van Kleeck (Eds.), *Children's language (Vol. 6, pp. 191–212)*. Hillsdale, NJ: Erlbaum.
- Gould, J. L. (1984). Natural history of honey bee learning. In P. Marler & H. S. Terrace (Eds.), *The biology of learning* (pp. 149–180). Berlin: Springer-Verlag.
- Gould, J. L., & Marler, P. (1984). Ethology and the natural history of learning. In P. Marler & H. S. Terrace (Eds.), *The biology of learning* (pp. 47–74). Berlin: Springer-Verlag.
- Hutchinson, J. E. (1984). *Constraints on children's implicit hypotheses about word meanings*. Unpublished doctoral dissertation, Stanford University, Stanford, CA.

- Hutchinson, J. E. (1986). Children's sensitivity to the contrastive use of object category terms. *Papers and Reports on Child Language Development*, 25, 49–55.
- Huttenlocher, J., & Smiley, P. (1987). Early word meanings: The case of object names. *Cognitive Psychology*, 19, 63–89.
- Immelmann, K. (1972). Sexual and other long-term aspects of imprinting in birds and other species. In D. Lehrman, R. Hinde, & E. Shaw (Eds.), *Advances in the study of behavior* (Vol. 4, pp. 147–174). New York: Academic Press.
- Karmiloff-Smith, A. (1979). *Language as a formal problem-space for children*. Paper presented at the MPG/NIAS Conference on "Beyond description in child language."
- Karmiloff-Smith, A., & Inhelder, B. (1975). If you want to get ahead, get a theory. *Cognition*, 3, 195–211.
- Keeton, W. (1974). The orientational and navigational basis of homing in birds. In D. Lehrman, J. Rosenblatt, R. Hinde, & E. Shaw (Eds.), *Advances in the study of behavior* (Vol. 5, pp. 48–132). New York: Academic Press.
- Keil, F. C. (1979). *Semantic and conceptual development: An ontological perspective*. Cambridge, MA: Harvard University Press.
- Keil, F. C. (1981). Constraints on knowledge and cognitive development. *Psychological Review*, 88, 197–227.
- Keil, F. C. (1990). Constraints on constraints: Surveying the epigenetic landscape. *Cognitive Science*, 14, 135–168.
- Leopold, W. P. (1939). *Speech development of a bilingual child: A linguist's record: Vol. 1. Vocabulary growth in the first two years*. Evanston, IL: Northwestern University Press.
- Leopold, W. P. (1949). *Speech development of a bilingual child: A linguist's record: Vol. 3. Grammar and general problems in the first two years*. Evanston, IL: Northwestern University Press.
- Liittschwager, J. C., & Markman, E. M. (1991). *Mutual exclusivity as a default assumption in second label learning*. Presentation at the Biennial Convention of the Society for Research in Child Development, Seattle, WA.
- Macnamara, J. (1982). *Names for things: A study of human learning*. Cambridge, MA: MIT Press.
- MacWhinney, B. (1989). *Making words make sense: Commentary on Merriman and Bowman*. Commentary on Merriman and Bowman's Monograph for the Society for Research in Child Development.
- Markman, E. M. (1987). How children constrain possible word meanings. In U. Neisser (Ed.), *The ecological and intellectual bases of categorization* (pp. 255–287). Cambridge, MA: Cambridge University Press.
- Markman, E. M. (1989). *Categorization in children: Problems of induction*. Cambridge, MA: MIT Press, Bradford Books.
- Markman, E. M. (1990). Constraints children place on word meanings. *Cognitive Science*, 14, 57–77.
- Markman, E. M. (in press). Ways in which children constrain word meanings. In E. Dromi (Ed.), S. Strauss (Series Ed.), *Language and cognition: A developmental perspective, The Fifth Annual Tel Aviv Workshop in Human Development*. Norwood, NJ: Ablex.
- Markman, E. M., & Callanan, M. A. (1983). An analysis of hierarchical classification. In R. Sternberg (Ed.), *Advances in the psychology of human intelligence*, Vol. 2, (pp. 325–365). Hillsdale, NJ: Erlbaum.
- Markman, E. M., & Hutchinson, J. (1984). Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations. *Cognitive Psychology*, 16, 1–27.
- Markman, E. M., & Wachtel, G. F. (1988). Children's use of mutual exclusivity to constrain the meanings of words. *Cognitive Psychology*, 20, 121–157.

- Marler, P., & Terrace, H. S. (1984). *The biology of learning*. Berlin: Springer-Verlag.
- Matthews, G. (1953). Navigation in the Manx Shearwater. *Journal of Experimental Biology*, 30, 370–396.
- Mayr, E. (1974). Behavior programs and evolutionary strategies. *American Scientist*, 62, 650–659.
- McShane, J. (1979). The development of naming. *Linguistics*, 17, 879–905.
- Merriman, W. E. (1986). Some reasons for the occurrence and eventual correction of children's naming errors. *Child Development*, 57, 942–952.
- Merriman, W. E. (1987). *Lexical contrast in toddlers: A reanalysis of the diary evidence*. Presentation at the Biennial Convention of the Society for Research in Child Development, Baltimore, MD.
- Merriman, W. E., & Bowman, L. L. (1989). The mutual exclusivity bias in children's word learning. *Monographs of the Society for Research in Child Development*, 54, Serial No. 220.
- Mervis, C. B. (1987). Child-basic object categories and early lexical development. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 201–233). Cambridge, MA: Cambridge University Press.
- Mervis, C. B. (1989). *Operating principles and early lexical development*. Paper presented at the Biennial Convention of the Society for Research in Child Development, Kansas City, MO.
- Nelson, K. (1973). Structure and strategy in learning to talk. *Monographs of the Society for Research in Child Development*, 38, Serial No. 149.
- Nelson, K. (1988). Constraints on word learning? *Cognitive Development*, 3, 221–246.
- Newport, E. L. (1990). Maturation constraints on language learning. *Cognitive Science*, 14, 11–28.
- Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.
- Quine, W. V. O. (1960). *Word and object*. Cambridge, MA: MIT Press.
- Rozin, P., & Kalat, J. W. (1971). Specific hungers and poison avoidance as adaptive specializations of learning. *Psychology Review*, 78, 459–486.
- Shettleworth, S. J. (1972). Constraints on learning. In D. Lehrman, R. Hinde, & E. Shaw (Eds.), *Advances in the study of behavior* (Vol. 4, pp. 1–68). New York: Academic Press.
- Shettleworth, S. J. (1984). Natural history and evolution of learning in nonhuman mammals. In P. Marler & H. S. Terrace (Eds.), *The biology of learning* (pp. 419–434). New York: Springer-Verlag.
- Soja, N., Carey, S., & Spelke, E. (1985). *Constraints on word learning*. Paper presented at the Biennial Convention of the Society for Research in Child Development, Toronto, Canada.
- Spelke, E. S. (1990). Principles of object perception. *Cognitive Science*, 14, 29–56.
- Waxman, S. R. (1990). Linguistic biases and the establishment of conceptual hierarchies: Evidence from preschool children. *Cognitive Development*, 5, 123–150.
- Waxman, S. R. (1989). Linking language and conceptual development: Linguistic cues and the construction of conceptual hierarchies. *Genetic Epistemologist*, 17, 13–19.
- Waxman, S. R., & Gelman, R. (1986). Preschooler's use of superordinate relations in classification. *Cognitive Development*, 1, 139–156.
- Waxman, S. R., & Kosowski, T. D. (1990). Nouns mark category relations: Toddlers' and preschoolers' word-learning biases. *Child Development*, 61, 1461–1473.

RECEIVED: February 15, 1990; REVISED: March 13, 1990